



(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention
of the grant of the patent:
18.11.1998 Bulletin 1998/47

(51) Int Cl.⁶: **G06T 7/00**

(21) Application number: **92301631.5**

(22) Date of filing: **26.02.1992**

(54) **Method and apparatus for optically determining the acceptability of products**

Verfahren und Gerät zur optischen Bestimmung der Annehmbarkeit von Produkten

Méthode et appareil de détermination optique de l'acceptabilité de produits

(84) Designated Contracting States:
**AT BE CH DE DK ES FR GB GR IT LI LU MC NL
PT SE**

(30) Priority: **27.02.1991 US 661809**

(43) Date of publication of application:
02.09.1992 Bulletin 1992/36

(73) Proprietor: **PHILIP MORRIS PRODUCTS INC.**
Richmond, Virginia 23261-6583 (US)

(72) Inventors:
• **Cox, Kenneth A.**
Midlothian, Virginia 23113 (US)

• **Dante, Henry M.**
Midlothian, Virginia 23112 (US)
• **Maher, Robert J.**
Midlothian, Virginia 23113 (US)

(74) Representative:
Lloyd, Patrick Alexander Desmond
Reddie & Grose
16 Theobalds Road
London WC1X 8PL (GB)

(56) References cited:
EP-A- 0 382 466 **US-A- 4 952 805**

Description

This invention relates to product inspection methods and apparatus, and more particularly to a method and apparatus for optically determining whether or not a product has an acceptable appearance.

For many products such as consumer goods like packaged foods, beverages, cleaning products, health and beauty aids, cigarettes, cigars, etc., it is very important that the external appearance of the product or its packaging be uniform and defect-free. Yet these products are typically produced in such large quantities and at such high speeds that some form of automated optical inspection is practically essential. It is highly desirable for an optical inspection system to be able to test all or substantially all parts of the product image so that defects of any kind occurring anywhere in the image can be detected. At the same time the inspection system should not reject products having minor but acceptable deviations from the ideal product.

For even a relatively simple product image such as a cigarette pack, an inspection system must be initially supplied with a tremendous amount of information in order to enable the system to inspect all or substantially all portions of the image with the sophistication required to discriminate between acceptable products (i.e., products having the ideal appearance or an appearance acceptably close to the ideal) and unacceptable products which should be rejected because of defects in appearance or appearance which is not sufficiently close to the ideal. Identifying and entering this information into the inspection system apparatus typically requires a very high level of skill and/or large amounts of operator time. Moreover, this data identification and entry task must be repeated each time a new or even slightly different product is to be inspected. Such an inspection system is revealed in EP-A-0 382 466.

The present invention aims to improve and simplify optical inspection systems.

Accordingly, there is provided a method of determining the acceptability of a product by generating a filter (F) from a first set of acceptable images (I) of the product, comparing the filter with each of a second set of images of the product to produce a processed value (P) for each image in the second set, the processed values having a distribution of values, comparing the filter with each of a third set of images of the product to produce a further processed value for each image in the third set, and comparing each further processed value to the distribution to determine the acceptability of the product associated with the image having the further processed value, characterized by:

generating from the distribution a first range (A) and a second range (B) of processed values, the first range comprising processed values associated with acceptable images of the product, and the second range being spaced from said first range and

comprising processed values associated with unacceptable images of the product;
determining whether each further processed value is outside both of the first and second ranges and, if so, selecting the associated image only if said image is acceptable; and
adaptively training the filter with the selected image to produce a modified filter whereby comparison of the modified filter with the selected image produces a modified processed value closer to the first range.

The invention also provides apparatus for determining the acceptability of a product by generating a filter (F) from a first set of acceptable images (I) of the product, comparing the filter with each of a second set of images of the product to produce a processed value (P) for each image in the second set, the processed values having a distribution of values, comparing the filter with each of a third set of images of the product to produce a further processed value for each image in the third set, and comparing each further processed value to the distribution to determine the acceptability of the product associated with the image having the further processed value, characterized by:

means for generating from the distribution a first range (A) and a second range (B) of processed values, the first range comprising processed values associated with acceptable images of the product, and the second range being spaced from the first range and comprising processed values associated with unacceptable images of the product;
means for determining whether each further processed value is outside both of the first and second ranges and, if so, selecting the associated image only if said image is acceptable; and
means for adaptively training the filter with the selected image to produce a modified filter whereby the comparison of the modified filter with the selected image produces a modified processed value closer to the first range.

Embodiments of the invention have the advantage that they may greatly reduce the level of operator skill and amount of operator time required to set up the system to inspect a new or different product.

In a preferred embodiment of the invention an optical inspection system forms an initial discriminant function or "filter" from a composite of a relatively small number of "first phase" images which the operator of the system determines to be acceptable images. A relatively simple technique (e.g., a logical OR function) is preferably used to form this composite. The ability to produce the initial discriminant function quickly using such a simple combination of a small number of images facilitates rapid "start-up" of the system. Thereafter the system uses the initial discriminant function to process a relatively large number of representative "second

phase" images in order to compute statistical information about the images in relation to the initial discriminant function. In particular, the system uses the initial discriminant function to compute a processed value for each second phase image. These processed values will typically have an approximately normal or Gaussian distribution. The upper and lower limits of a central portion of this distribution containing a first statistically large number of the processed values are identified as first threshold values. The upper and lower limits of a central portion of this distribution containing a second statistically even larger number of the processed values are identified as second threshold values.

In a subsequent third phase of the operation of the system, the first and second threshold values are used in the processing of a relatively large number of "third phase" images. In particular, the system uses a discriminant function (initially the above-mentioned initial discriminant function) to compute a processed value for each successive third phase image. If this processed value for a given third phase image is between the first threshold values, that third phase image is automatically used to refine (e.g., using a Widrow-Hoff-type adaptive training process) the discriminant function for subsequent use. If the processed value for a third phase image is not between the second threshold values, that third phase image is automatically discarded. As a third possibility, if the processed value for a third phase image is not between the first threshold values but is between the second threshold values, the operator of the system is given the choice as to whether or not that image should be discarded (i.e., because the image looks unacceptable) or used to refine the discriminant function for subsequent use (i.e., because the image looks acceptable).

When the third phase is completed, the system is ready for actual product inspection using the refined discriminant function and the first threshold values. In actual product inspection, the system uses the refined discriminant function to compute a processed value for each product image. If the processed value for a product image is between the first threshold values, the product is accepted as having an acceptable appearance. If the processed value for a product image is not between the first threshold values, the product is rejected as having an unacceptable appearance.

The system greatly reduces the level of operator skill and amount of operator time required to set the system up for a product inspection task. The operator is only required to identify a relatively small number of acceptable images (e.g., 25) during the first phase. The initial discriminant function is then computed automatically using a simple and rapid technique such as a logical OR of the small number of first phase images. The entire second phase may also be automatic. And during the third phase, the operator is only required to decide on the acceptability of the relatively small number of images whose processed values fall outside the first threshold values but between the second threshold values.

An embodiment of the invention will now be described by way of example, and with reference to the accompanying drawings in which:

FIG. 1 is a simplified schematic block diagram of an illustrative embodiment of an optical product inspection apparatus embodying the invention; FIGS. 2a and 2b (referred to collectively as FIG. 2) are a flow chart of an illustrative optical product inspection method embodying the invention; FIGS. 3a-3c (referred to collectively as FIG. 3) are a flow chart showing in greater detail, the first training phase shown in FIG. 2; FIGS. 4a and 4b (referred to collectively as FIG. 4) are a flow chart showing in greater detail, the second training phase shown in FIG. 2; FIGS. 5a-5c (referred to collectively as FIG. 5) are a flow chart showing in greater detail the third training phase shown in FIG. 2; FIG. 6 is a flow chart showing in greater detail the product inspection steps shown in FIG. 2; FIG. 7 is a diagram of a dot product spectrum which is useful in explaining certain features of the embodiment described; FIG. 8 is a histogram diagram useful in explaining certain features of the embodiment described; and FIG. 9 shows several equations which may be employed in the embodiment of the invention.

As shown in FIG. 1, a typical product inspection system 10 constructed in accordance with this invention includes conveyor apparatus 20 for conveying the objects or products 12 to be inspected, one after another, from left to right as viewed in the FIG. Each time conventional product sensor 22 detects a product 12 at a predetermined location A opposite conventional camera 24, conventional processor 26 (which includes conventional imaging hardware) causes conventional light sources 30 to briefly illuminate the product, thereby allowing camera 24 to capture what is effectively a still image of the product. This still image is fed to processor 26 which digitizes and further processes the image. Processor 26 is augmented by conventional video display 32 and conventional data entry device 34 (e.g., a keyboard, mouse, and/or touch screen elements associated with display 32). Processor 26 can cause display 32 to display a product image captured by camera 24, and can augment that display with other information such as the outline of an acceptable product image and/or outlines of certain features of an acceptable product image. The operator may use this augmenting information to help determine whether the product image being displayed is acceptable. The operator may use data entry device 34 to control the overall operation of the system, as well as to respond to inquiries from the system (e.g., as to whether or not the operator judges the product image currently shown on display 32 to be acceptable).

The system may be set up to perform a product in-

spection by operating it substantially as though it were inspecting products, i.e., by using conveyor 20 to convey representative products one after another past camera 24 and by using the other elements of the system to process the images of those products as described in detail below. During actual product inspection, processor 26 determines whether the image of each successive product 12 is acceptable, and when that product reaches a controllable branch 20B in conveyor 20, processor 26 controls that branch so that acceptable products 12A are directed to accepted product conveyor 20A, while unacceptable products 12R are directed to rejected product conveyor 20R.

While FIG. 1 suggests that system 10 operates on a single elevational image of products 12, it will be apparent to those skilled in the art that the system could be set up to test multiple images of the products taken from different angles and including perspective views so that as many surfaces of the objects are inspected as are desired. Similarly, although the system will be explained in terms of monochrome (e.g., black and white) images, it will be apparent to those skilled in the art how the system can be modified to inspect in full color. Thus camera 24 may be a conventional NTSC or RGB compatible camera. Processor 26 may be a suitably programmed conventional 386 personal computer workstation such as a CAT386 workstation available from Comark Corp. of Medfield, Massachusetts with a conventional IM-1280 imaging hardware system available from MATROX Electronic Systems Limited of Dorval, Quebec, Canada.

An overview of a preferred embodiment of the method of this invention is shown in FIG. 2. Basically the depicted embodiment comprises a training portion, including three successive phases 1, 2, and 3 (shown in boxes 100, 200, and 300, respectively), and actual product inspection (shown in box 400). During the three training phases, the system "learns", by appropriately processing product images with appropriate but relatively limited input from the human operator of the system, how to discriminate between good and bad images. Thereafter, during actual product inspection, the system uses this "knowledge" to accept or reject products.

In training phase 1 (step 100 in FIG. 2) an initial discriminant function F (which may be thought of as a two-dimensional matrix commensurate with the two-dimensional data for the product images or product image portions to be inspected) is formed from the data I for a relatively small number of "first phase" images. Although this initial discriminant function could be computed in many other ways in accordance with this invention, in the preferred embodiment (shown in detail in FIG. 3) a relatively simple technique (i.e., a logical OR of the phase 1 images) is used in order to allow a relatively small and inexpensive processor 26 to perform the necessary calculations without requiring more time than the operator of the system needs to provide the necessary inputs regarding each successive first phase image. Ac-

cordingly, as shown in FIG. 3 training phase 1 starts with step 102, and in step 104 various program constants are initialized and inputs are read (e.g., from the memory which is part of processor 26 and/or from data entry device 34). For example, step 104 may include selection of an image outline overlay to be displayed with product images on display 32 to help the operator judge the acceptability of images. Step 104 may also include selection of the boundaries of the portion or portions of the image to be processed. As another example, step 104 may include selection of a threshold to be used in binarizing the image data as discussed below. Any other necessary system initialization tasks may be performed as part of step 104.

In step 106 the system acquires the data for the first of the first phase images. This is done by having camera 24 capture a product image as described above. Processor 26 then digitizes this image in full gray scale and causes display 32 to display this gray scale image with any augmenting information (such as an outline overlay) selected in step 104. The operator then indicates (via data entry device 34) whether or not the displayed image is acceptable. If so, control passes to step 108. If not, step 106 is repeated with new product images until an image acceptable to the operator is found.

In step 108 the first acceptable image is preprocessed. This preferably includes edge detecting the gray scale image so that pixels at or near significant changes in image brightness are emphasized (e.g., increased in value) relative to other pixels which are de-emphasized (e.g., decreased in value). Edge detection is a well-known technique which is discussed in more detail, for example, in European Patent Application EP-A-0382466 claiming priority from U.S. application serial number. 308,739, filed February 9, 1989.

After edge detection, the edge detected image is preferably binarized so that all pixels having values on one side of a predetermined binarization threshold level (which may have been selected in step 104) are assigned one binary value (e.g., 1), while all pixels having values on the other side of the binarization threshold level are assigned the other binary value (e.g., 0).

In step 110 the initial discriminant function F is set equal to the first acceptable image data from step 108.

In step 112 the sum of the dot products of the discriminant function and the phase 1 image data is initialized. Because at this point F and I are the same, the initial dot product of F and I is just the sum of the pixel values of I .

In step 114 an index value i is set equal to 1.

In step 116 the system acquires the next acceptable image. Step 116 is therefore an exact repetition of above-described step 106.

In step 118 the data for the next acceptable image (acquired in step 116) is preprocessed exactly as described above in connection with step 108.

In step 120 the initial discriminant function is updated with the new image data by computing the logical OR

of the new image data and the old initial discriminant function data to produce a new initial discriminant function. In other words, for each pixel location in which either or both of the image data and the old initial discriminant function data are 1, the new initial discriminant data value is 1, while for each pixel location in which both the image data and the old initial discriminant data are 0, the new initial discriminant function data value is 0.

In step 122 the sum of the dot products of the discriminant function and the phase 1 image data is updated for the current image. Because the 1-valued pixel locations in each image are always a subset of the 1-valued pixel locations in F, each new dot product is just the sum of the pixel values in the current image I.

In step 124 the index i is incremented by 1, and in step 126 the new value of i is tested to determine whether it is greater than or equal to 25. This is an arbitrary number which determines how many first phase images will be used to compute the initial discriminant function. Although any other number could be used, 25 has been found to give good results. If i has not yet reached 25, control returns to step 116 and steps 116-126 are repeated until the test in step 126 is satisfied and control consequently passes to step 128.

In step 128 the average of the dot products of F and each of the first phase images is computed by dividing PSUM by 25 (the number of first phase images).

In step 130 training phase 2 (step 200 in FIG. 2) begins. The initial discriminant function F from the last performance of step 120 and the average dot product are saved.

In step 200 the initial discriminant function F is used to compute statistical information about the images being processed. Again, although this can be done in other ways in accordance with this invention, a preferred embodiment of step 200 is shown in FIG. 4 and will now be described by way of illustration.

Training phase 2 starts in step 202. In step 204 index value i is set equal to 1, variable SUM is set equal to 0, and variable SUMSQRS (for sum of squares) is also set equal to 0.

In step 205 the initial binary discriminant function F is converted to bipolar form using initial positive and negative values such that the final discriminant function values take advantage of the full arithmetic range of processor 26. To reflect this in the average dot product, the average dot product is also scaled by the same scale factor in step 205. For example, if processor 26 performs 8-bit arithmetic with values between -128 and +127, the initial values now used in function F may be -50 (for pixel locations where the F value was formerly 0) and +50 (for pixel values where the F value was formerly 1), and the average dot product from step 128 may be multiplied by 50.

In step 206 a product image is acquired in a manner similar to above-described step 106, except that during training phase 2 the operator of the system is not re-

quired to determine whether the image is acceptable. Accordingly, all the images received during phase 2 are used. These images can therefore be expected to exhibit the normal range of variation for the product images that the system will subsequently encounter during actual product inspection. In addition, because no interaction with the operator of the system is required during this phase, the phase 2 images can be processed much faster (e.g., at actual product inspection rates) than the phase 1 images.

In step 208 the image data acquired in step 206 is preprocessed exactly as described above in connection with step 108.

In step 210 the dot product P of the rescaled initial discriminant function F from step 205 and the image data I from step 208 is calculated. This calculation involves multiplying the value of F at each pixel location by the value of I at that pixel location, and then summing all of the resulting products to produce the dot product P. Elsewhere in this specification the more generic term "processed value" is sometimes used for the dot product P. It will be noted that if I is identical to F, P will be a certain number, but if I differs from F at certain pixel locations, P will be greater or less than that number. The amount by which P differs from that number is a measure of how much I differs from F. In practice, the values of P will typically exhibit an approximately normal (i.e., approximately Gaussian) distribution about some mean or average value.

In step 212 the variable SUMSQRS is incremented by the square of the value of P from step 210, and the index value i is incremented by 1.

In step 214 the index value i is compared to an arbitrary number (e.g., 1000) which is the predetermined number of images to be processed in phase 2. Although any sufficiently large number of images can be processed in phase 2, 1000 images have been found to give good results. If the test in step 214 is not satisfied, control returns to step 206 where processing of the next phase 2 image begins. When 1000 phase 2 images have been processed as described above, the test in step 214 is satisfied and control passes from step 214 to step 216.

In step 216 the index value i is reduced by 1 to reverse the last incrementing of that value.

In step 218 the rescaled average dot product from step 205 and the value of the SUMSQRS variable are used to compute the standard deviation of the previously computed dot products.

In step 220 two first threshold values and two second threshold values are selected so that the distribution of phase 2 dot products is subdivided by these threshold values as shown in FIG. 7. For example, the first threshold values may be chosen so that a fraction f1 of the dot products are greater than the upper one of these threshold values and the same fraction of dot products are less than the lower one of these threshold values. The fraction f1 is preferably significantly greater than one-half

the fraction of images which are expected to be defective in order to minimize the possibility that any unacceptable images have dot products that are greater than the upper or less than the lower of these first threshold values. Images with dot products between the first threshold values are therefore automatically acceptable as indicated in FIG. 7.

The second threshold values are chosen so that a smaller fraction f_2 of the dot products are greater than the upper one of these threshold values and the same smaller fraction of dot products are less than the lower one of these threshold values. The fraction f_2 is preferably significantly smaller than one-half the fraction of images which are expected to be defective in order to minimize the possibility that any acceptable images have dot products that are less than the lower or greater than the upper one of these second threshold values. Images with dot products outside the region bounded by the second threshold values are therefore automatically rejectable as "gross rejects" as indicated in FIG. 7. Images with dot products outside the region bounded by the first threshold values but inside the region bounded by the second threshold values are "marginal rejects" as indicated in FIG. 7. Operator intervention is required to determine whether such an image should be accepted or rejected.

It may be convenient and appropriate to choose the threshold values described above assuming the distribution of dot products to be Gaussian as shown, for example, in FIG. 8, and therefore characterized by a standard deviation (given by the equation in step 128). In that case, the thresholds can be defined by the equations shown in FIG. 9. The average dot product used in these equations is the rescaled average dot product from step 205. The alpha coefficients used in these equations with the standard deviation are selected so as to achieve the target fractions f_1 and f_2 for a Gaussian distribution. These values can be readily selected with the aid of available tables of the properties of the Gaussian distribution. The most preferred approach is to select the first threshold values without assuming a Gaussian distribution (i.e., as described prior to the above discussion of the Gaussian distribution), and to use the second method (i.e., the Gaussian distribution assumption) to select the second threshold values. Note that in FIG. 8 the region A corresponds to the "acceptable" region of FIG. 7, the regions B correspond to the "gross reject" regions of FIG. 7, and the regions C correspond to the "marginal reject" regions of FIG. 7. Thus region A includes dot products known to be associated with clearly acceptable images, whereas regions B include dot products known to be associated with clearly unacceptable images. Regions C are those along the distribution of dot products P which may be marginally acceptable. Adaptive training is performed in phase 3 as discussed below with respect to dot products lying in region A, and also with respect to dot products lying in regions C which the operator of the system determines to be acceptable.

After the second threshold values are calculated in step 220, control passes to step 222 to begin training phase 3 (step 300 in FIG. 2). The first and second threshold values from step 220 are saved, as are F and the rescaled average dot product from step 205.

In training phase 3 (step 300 in FIG. 2) the statistical information from phase 2 is used with the image data from another statistically significant number of images to refine the initial discriminant function F . Again, although this can be done in other ways in accordance with this invention, a preferred embodiment of training phase 3 is shown in FIG. 5 which will now be described by way of illustration.

As shown in FIG. 5, training phase 3 begins in step 302, and in step 310 an index value i is initialized to 0, and a counter -- used during phase 3 to count the number of marginally acceptable images which the operator of the system decides to accept -- is also initialized to 0.

In step 312 a phase 3 image is acquired exactly as described above in connection with step 206, and in step 314 the data for this image is preprocessed as described above in connection with step 208.

In step 316 the dot product P of the discriminant function and the image data from step 314 is calculated.

In step 318 the dot product value P from step 316 is compared to the second threshold values from step 220. If P is outside the range bounded by these second threshold values, control passes to step 320 where the image is rejected and control is then returned to step 312 to begin the acquisition and processing of the next phase 3 image. On the other hand, if P is in the range bounded by the second threshold values, control passes from step 318 to step 322.

In step 322 the value of P from step 316 is compared to the first threshold values from step 220. If P is outside the range bounded by these first threshold values, control passes to step 324. Step 324 is reached whenever the image is neither automatically rejectable as unacceptable (because the associated dot product is outside the limits defined by the second threshold values), nor automatically acceptable (because the associated dot product is inside the limits defined by the first threshold values). Accordingly, in step 324 the operator of the system is asked to intervene and decide whether or not the image is acceptable. The image is displayed on display 32 (as in connection with step 106 above). If the operator responds (again as in connection with step 106) that the image is unacceptable, control passes to step 320 where the image is rejected, and thereafter processing of the next phase 3 image begins as described above. On the other hand, if the operator responds that the image is acceptable, control passes from step 324 to step 326 where the counter $nacc$ is incremented. Thereafter, control passes to step 328. Returning to the other branch from step 322, if P is not outside the limits defined by the first threshold values, the image is automatically acceptable and control passes directly from step

322 to step 328.

Step 328 is performed only when the current third phase image has been determined to be an acceptable image. In most cases the system will have made this determination automatically because the dot product P for the image is between the first threshold values and the image is therefore obviously acceptable. In a few cases, however, the operator will have been required to assist with this determination as described above in connection with step 324. Accordingly, for the most part the processing of images can proceed as rapidly during phase 3 as during phase 2. Only rarely will the operator be required to intervene as a result of the performance of step 324. Moreover, operator intervention should be required even less frequently as phase 3 proceeds and the discriminant function is progressively refined as will now be described.

Step 328 begins the process of refining the rescaled discriminant function using the data from the image which has just been determined to be acceptable. This discriminant function refining process is repeated for each acceptable phase 3 image. In step 328 an "error" value equal to the difference between the average P value from step 205 and P from step 316 is calculated. In step 329 a value N equal to the number of pixels which are "on" in the image data is calculated. In step 330 a correction value equal to the error value from step 328 divided by the value of N from step 329 is calculated. In step 332 the binary image data for the current phase 3 image is rescaled using the correction value from step 330. In particular, each pixel value of 1 is changed to the correction value, while each pixel value of 0 is unaltered.

In step 334 the rescaled discriminant function is refined by incrementing each pixel value by the value associated with that pixel in the rescaled image data from step 332. Step 334 is an "adaptive training" step analogous to the Widrow-Hoff training algorithm sometimes used in signal processing (see, for example, B. Widrow and S.D. Stearns, Adaptive Signal Processing, Prentice-Hall, Englewood Cliffs, 1985). Accordingly, as step 334 is performed for successive acceptable third phase images, the rescaled discriminant function becomes better and better at producing dot products (as in step 316) which are clearly differentiated between those associated with acceptable images (P within the range bounded by the first threshold values) and those associated with unacceptable images (P outside the range bounded by the second threshold values). Accordingly, as phase 3 progresses, there should be less and less need to perform step 324, and the amount of input required from the operator of the system should decrease.

In step 336 the index value i is incremented. In step 338 this index value is compared to a phase 3 cut-off value (e.g., 2000 acceptable phase 3 images). If the index value is less than the cut-off value, control passes from step 338 to step 312 where processing of the next phase 3 image begins. As soon as step 338 detects that the index value has reached the cut-off value, control

passes from step 338 to step 340.

In step 340 the ratio of the counter value nacc to the index value i is compared to a predetermined threshold value. If this ratio exceeds the threshold value, the system is still tentatively rejecting too many images which the operator of the system has found acceptable in step 324. This indicates that the discriminant function F' is still in need of further refinement. Accordingly, control passes from step 340 to step 310 where the processing of another series of phase 3 images begins again. On the other hand, if the ratio in step 340 is less than the threshold, the refining of discriminant function F' is judged complete, and training phase 3 is concluded by passing control to step 342 where actual product inspection begins (step 400 in FIG. 2).

An illustrative embodiment of actual product inspection (step 400 in FIG. 2) is shown in FIG. 6. This process begins with step 402, and in step 404 an image is acquired as in step 206. In step 406 the data for this image is preprocessed as in step 208. In step 408 the dot product P of the refined discriminant function from training phase 3 and the image data from step 406 is calculated. In step 410 P is tested to determine whether or not it is in the range between the first threshold values from step 220. If so, the system deems the image acceptable and control passes to step 412 in order to accept the product (i.e., direct it to accepted product conveyor 20A in FIG. 1). If the step 410 test is not satisfied, the system deems the image unacceptable and control passes to step 414 in order to reject the product (i.e., direct it to rejected product conveyor 20R in FIG. 1). After either step 412 or 414, control returns to step 404 to begin processing of the next image.

It will be understood that the foregoing is merely illustrative of the principles of this invention and that various modifications can be made by those skilled in the art without departing from the scope of the invention. For example, each image formed by elements 24 and 26 can be broken down into a plurality of predetermined segments, and the data for each of the segments can be separately processed just as the data for the whole image is processed in the foregoing discussion. During actual product inspection all segments must satisfy the test of step 410 in order for the product to be accepted. As another example of a modification within the scope of this invention, discriminant function F' can continue to be refined during actual product inspection by updating it in accordance with any acceptable product image or images exactly as described above in connection with training phase 3. In another modification within the scope of the invention bipolar values -1 and +1 are used instead of binary values throughout and the discussion of, for example, steps 108, 118, 208, 314 and 406 should be construed accordingly. It should be understood that both binary and bipolar values are merely examples and any other two values can be used instead and are within the scope of the invention.

Claims

1. A method of determining the acceptability of a product (12) by generating a filter (F) from a first set of acceptable images (I) of the product, comparing the filter with each of a second set of images of the product to produce a processed value (P) for each image in the second set, the processed values having a distribution of values, comparing the filter with each of a third set of images of the product to produce a further processed value for each image in the third set, and comparing each further processed value to the distribution to determine the acceptability of the product associated with the image having the further processed value, characterized by:

generating (220) from the distribution a first range (A) and a second range (B) of processed values, the first range comprising processed values associated with acceptable images of the product, and the second range being spaced from said first range and comprising processed values associated with unacceptable images of the product;
determining (318, 322, 324) whether each further processed value is outside both of the first and second ranges and, if so, selecting the associated image only if said image is acceptable; and
adaptively training (334) the filter with the selected image to produce a modified filter whereby comparison of the modified filter with the selected image produces a modified processed value closer to the first range.

2. A method according to Claim 1, characterized in generating the filter from the first set of acceptable images of the product comprises combining (120) the first set of acceptable images using a logical OR operation.

3. A method as claimed in Claim 2, further characterized in that the step of combining the first set of acceptable images includes the steps of:

edge detecting (108) each image; and
associating (108) a first value with each portion of each edge detected image which has a value on one side of a predetermined threshold value, and associating a second value with all other portions of each edge detected image prior to combining the images using the logical OR operation.

4. A method according to Claim 3, characterized in that the first and second values are binary 0 and 1.

5. A method according to any of Claims 1 to 4, char-

acterized in that the adaptive training of the filter is analogous to a Widrow-Hoff training algorithm.

6. A method according to any of Claims 1 to 5, characterized in that the selected image is selected manually (324) based on the appearance of the product associated with the image.

7. A method according to any of Claims 1 to 6, characterized by adaptively training the filter with the associated image if the further processed value is in the first range.

8. A method according to any preceding claim, wherein the processed values of each image in the second set are the dot product of the filter and each of the second set of images.

9. A method according to any preceding claim, characterized by sub-dividing each of the acceptable images into a plurality of image portions and performing all of the foregoing steps upon at least some of the image portions individually.

10. Apparatus for determining the acceptability of a product (12) by generating a filter (F) from a first set of acceptable images (I) of the product, comparing the filter with each of a second set of images of the product to produce a processed value (P) for each image in the second set, the processed values having a distribution of values, comparing the filter with each of a third set of images of the product to produce a further processed value for each image in the third set, and comparing each further processed value to the distribution to determine the acceptability of the product associated with the image having the further processed value, characterized by:

means (26, 220) for generating from the distribution a first range (A) and a second range (B) of processed values, the first range comprising processed values associated with acceptable images of the product, and the second range being spaced from the first range and comprising processed values associated with unacceptable images of the product;
means (26, 318, 322, 324) for determining whether each further processed value is outside both of the first and second ranges and, if so, selecting the associated image only if said image is acceptable; and
means (26, 334) for adaptively training the filter with the selected image to produce a modified filter whereby the comparison of the modified filter with the selected image produces a modified processed value closer to the first range.

11. Apparatus according to Claim 10, characterized in

that the filter is generated from the first set of acceptable images of the product by means (26, 120) for combining the first set of acceptable images using a logical OR operation.

12. Apparatus according to Claim 11, characterised in that the means for combining the first set of acceptable images comprises:

means (26, 108) for edge detecting each image; and

means (26, 108) for associating a first value with each portion of each edge detected image which has a value on one side of a predetermined threshold value, and associated a second value with all other portions of each edge detected image prior to combining the images using the logical OR operation.

13. Apparatus according to Claim 12, characterized in that the first and second values are binary 1 and 0.

14. The apparatus as claimed in any of Claims 10 to 13, characterized in that the means (26, 334) for adaptively training performs a function analogous to a Widrow-Hoff training algorithm.

15. Apparatus according to any of Claims 10 to 14, characterized in that the means for determining includes means (26, 324) for allowing manual selection based on the appearance of the product associated with the image.

16. Apparatus according to any of Claims 10 to 15, characterized by means (26) for adaptively training the filter with an image associated with a further processed value which is in the first range.

17. Apparatus according to any of Claims 10 to 16, characterized by means for forming the dot product of the filter and each of the second set of images to form the processed values of each image in the said second set.

18. Apparatus according to any of Claims 10 to 17, characterized by a video camera (24) for forming at least one of the first, second and third sets of images, and by means for positioning a product in the field of view of the video camera.

19. Apparatus according to Claim 18, characterized by means for illuminating the product in the field of view of the video camera.

20. Apparatus according to Claim 19 or 20, characterized in that the positioning means comprises a conveyor (20) for conveying products are after another through the field of view of the video camera.

Patentansprüche

1. Verfahren zur Bestimmung der Annehmbarkeit eines Produkts (12) durch Erzeugen eines Filters (F) aus einem ersten Satz annehmbarer Bilder (I) des Produkts, Vergleichen des Filters mit jedem Bild eines zweiten Satzes Bilder des Produkts, um einen verarbeiteten Wert (P) für jedes Bild in dem zweiten Satz zu erzeugen, wobei die verarbeiteten Werte eine Werteverteilung besitzen, Vergleichen des Filters mit jedem Bild eines dritten Satzes Bilder des Produkts, um einen weiterverarbeiteten Wert für jedes Bild in dem dritten Satz zu erzeugen, und Vergleichen eines jeden weiterverarbeiteten Wertes mit der Verteilung, um die Annehmbarkeit des Produkts, das zu dem Bild gehört, welches den weiterverarbeiteten Wert hat, zu bestimmen, gekennzeichnet durch:

Erstellen (220) eines ersten Bereichs (A) und eines zweiten Bereichs (B) verarbeiteter Werte aus der Verteilung, wobei der erste Bereich verarbeitete Werte umfaßt, die zu annehmbaren Bildern des Produkts gehören, und der zweite Bereich getrennt vom ersten Bereich ist und verarbeitete Werte umfaßt, die zu unannehmbaren Bildern des Produkts gehören, Bestimmen (318, 322, 324), ob jeder weiterverarbeitete Wert sowohl außerhalb des ersten als auch außerhalb des zweiten Bereichs liegt, und, falls dies der Fall ist, Auswählen des zugehörigen Bildes nur, wenn dieses annehmbar ist, und

Adaptives Trainieren (334) des Filters mit dem ausgewählten Bild, um einen modifizierten Filter zu erzeugen, wodurch der Vergleich des modifizierten Filters mit dem ausgewählten Bild einen modifizierten verarbeiteten Wert erzeugt, der näher am ersten Bereich liegt.

2. Verfahren nach Anspruch 1, gekennzeichnet durch Erzeugen des Filters aus dem ersten Satz annehmbarer Bilder des Produkts, indem der erste Satz annehmbarer Bilder durch Anwendung einer logischen ODER-Verknüpfung kombiniert wird (120).

3. Verfahren nach Anspruch 2, ferner dadurch gekennzeichnet, daß der Schritt der Kombination des ersten Satzes annehmbarer Bilder die Schritte umfaßt:

Detektieren (108) des Randes eines jeden Bildes und

Zuordnen (108) eines ersten Wertes zu jedem Teil eines jeden randdetektierten Bildes, das einen Wert auf einer Seite eines vorher festgesetzten Schwellenwertes hat, und Zuordnen eines zweiten Wertes zu allen anderen Teilen ei-

nes jeden randdetektierten Bildes vor der Kombination der Bilder durch Anwenden der logischen ODER-Verknüpfung.

4. Verfahren nach Anspruch 3, dadurch gekennzeichnet, daß die ersten und zweiten Werte binär 0 und 1 sind. 5
5. Verfahren nach irgendeinem der Ansprüche 1 bis 4, dadurch gekennzeichnet, daß das adaptive Training des Filters analog zu einem Widrow-Hoff-Trainingsalgorithmus ist. 10
6. Verfahren nach irgendeinem der Ansprüche 1 bis 5, dadurch gekennzeichnet, daß das ausgewählte Bild manuell ausgewählt wird (324), basierend auf dem Aussehen des zu dem Bild gehörigen Produkts. 15
7. Verfahren nach irgendeinem der Ansprüche 1 bis 6, gekennzeichnet durch adaptives Trainieren des Filters mit dem zugehörigen Bild, wenn der weiterverarbeitete Wert im ersten Bereich liegt. 20
8. Verfahren nach irgendeinem vorhergehenden Anspruch, bei dem die verarbeiteten Werte eines jeden Bildes im zweiten Satz das Skalarprodukt des Filters und eines jeden Bildes des zweiten Satzes sind. 25
9. Verfahren nach irgendeinem vorhergehenden Anspruch, gekennzeichnet durch Unterteilen eines jeden annehmbaren Bildes in eine Vielzahl von Bildteilen und einzelnes Anwenden sämtlicher vorhergehender Schritte auf wenigstens einige der Bildteile. 30
10. Vorrichtung zur Bestimmung der Annehmbarkeit eines Produkts (12) durch Erzeugen eines Filters (F) aus einem ersten Satz annehmbarer Bilder (I) des Produkts, Vergleichen des Filters mit jedem Bild eines zweiten Satzes Bilder des Produkts, um einen verarbeiteten Wert (P) für jedes Bild in dem zweiten Satz zu erzeugen, wobei die verarbeiteten Werte eine Werteverteilung besitzen, Vergleichen des Filters mit jedem Bild eines dritten Satzes Bilder des Produkts, um einen weiterverarbeiteten Wert für jedes Bild in dem dritten Satz zu erzeugen, und Vergleichen eines jeden weiterverarbeiteten Wertes mit der Verteilung, um die Annehmbarkeit des Produkts, das zu dem Bild gehört, welches den weiterverarbeiteten Wert hat, zu bestimmen, gekennzeichnet durch: 35

Mittel (26, 220), um aus der Verteilung einen ersten Bereich (A) und einen zweiten Bereich (B) verarbeiteter Werte zu erzeugen, wobei der erste Bereich verarbeitete Werte umfaßt, die zu 55

annehmbaren Bildern des Produkts gehören, und der zweite Bereich getrennt vom ersten Bereich ist und verarbeitete Werte umfaßt, die zu unannehmbaren Bildern des Produkts gehören,

Mittel (26, 318, 322, 324), um zu bestimmen, ob jeder weiterverarbeitete Wert sowohl außerhalb des ersten als auch außerhalb des zweiten Bereichs liegt, und, falls dies der Fall ist, um das zugehörige Bild nur auszuwählen, wenn dieses annehmbar ist, und

Mittel (26, 334) zum adaptiven Training des Filters mit dem ausgewählten Bild, um einen modifizierten Filter zu erzeugen, wodurch der Vergleich des modifizierten Filters mit dem ausgewählten Bild einen modifizierten verarbeiteten Wert erzeugt, der näher am ersten Bereich liegt.

11. Vorrichtung gemäß Anspruch 10, dadurch gekennzeichnet, daß der Filter aus dem ersten Satz annehmbarer Bilder des Produkts durch Mittel (26, 120) zur Kombination des ersten Satzes annehmbarer Bilder durch Anwendung einer logischen ODER-Verknüpfung erzeugt wird. 20

12. Vorrichtung gemäß Anspruch 11, dadurch gekennzeichnet, daß das Mittel zur Kombination des ersten Satzes annehmbarer Bilder umfaßt: 25

Mittel (26, 108) zum Detektieren des Randes eines jeden Bildes und

Mittel (26, 108) zur Zuordnung eines ersten Wertes zu jedem Teil eines jeden randdetektierten Bildes, das einen Wert auf einer Seite eines vorher festgesetzten Schwellenwertes hat, und zur Zuordnung eines zweiten Wertes zu allen anderen Teilen eines jeden randdetektierten Bildes vor der Kombination der Bilder durch Anwenden der logischen ODER-Verknüpfung. 30

13. Vorrichtung gemäß Anspruch 12, dadurch gekennzeichnet, daß die ersten und zweiten Werte binär 1 und 0 sind. 35

14. Vorrichtung gemäß irgendeinem der Ansprüche 10 bis 13, dadurch gekennzeichnet, daß das Mittel (26, 334) zum adaptiven Training des Filters eine zu einem Widrow-Hoff-Trainingsalgorithmus analoge Funktion ausübt. 40

15. Vorrichtung gemäß irgendeinem der Ansprüche 10 bis 14, dadurch gekennzeichnet, daß das Mittel zur Bestimmung Mittel (26, 324) zur Ermöglichung einer manuellen Auswahl, basierend auf dem Aussehen des zu dem Bild gehörigen Produkts, umfaßt. 45

16. Vorrichtung gemäß irgendeinem der Ansprüche 10 bis 15, gekennzeichnet durch ein Mittel (26) zum adaptiven Training des Filters mit einem Bild, das zu einem weiterverarbeiteten Wert gehört, der im ersten Bereich liegt. 5
17. Vorrichtung gemäß irgendeinem der Ansprüche 10 bis 16, gekennzeichnet durch Mittel zur Bildung des Skalarprodukts des Filters und eines jeden Bildes des zweiten Satzes, um die weiterverarbeiteten Werte von jedem Bild in dem zweiten Satz zu bilden. 10
18. Vorrichtung gemäß irgendeinem der Ansprüche 10 bis 17, gekennzeichnet durch eine Videokamera (24) zur Bildung von wenigstens einem Bild des ersten, zweiten und dritten Satzes und ein Mittel, um ein Produkt in das Sichtfeld der Videokamera zu positionieren. 15
19. Vorrichtung gemäß Anspruch 18, gekennzeichnet durch Mittel zur Ausleuchtung des Produkts im Sichtfeld der Videokamera. 20
20. Vorrichtung gemäß Anspruch 19 oder 20, dadurch gekennzeichnet, daß das Positionierungsmittel ein Fördermittel (20) zur Beförderung der Produkte nacheinander durch das Sichtfeld der Videokamera umfaßt. 25

Revendications

1. Procédé permettant de déterminer l'acceptabilité d'un produit (12) en générant un filtre (F) à partir d'un premier ensemble d'images acceptables (I) du produit, en comparant le filtre à chacune d'un deuxième ensemble d'images du produit pour produire une valeur traitée (P) pour chaque image du deuxième ensemble, les valeurs traitées ayant une distribution de valeurs, en comparant le filtre à chacune d'un troisième ensemble d'images du produit pour produire une autre valeur traitée pour chaque image du troisième ensemble, et en comparant chaque autre valeur traitée à la distribution pour déterminer l'acceptabilité du produit associé à l'image ayant l'autre valeur traitée, caractérisé en ce que le procédé comprend les étapes consistant : 35
- à générer (220) à partir de la distribution une première plage (A) et une seconde plage (B) de valeurs traitées, la première plage comprenant des valeurs traitées associées à des images acceptables du produit et la seconde plage étant espacée de ladite première plage et comprenant des valeurs traitées associées à des images inacceptables du produit; 40
- à déterminer (318, 322, 324) si chaque autre valeur traitée est en dehors à la fois de la pre- 45
- 50
- 55

mière et de la seconde plages et, si c'est le cas, à sélectionner l'image associée uniquement si ladite image est acceptable; et à entraîner de manière adaptative (334) le filtre avec l'image sélectionnée pour produire un filtre modifié de telle sorte que la comparaison du filtre modifié à l'image sélectionnée produise une valeur traitée modifiée plus proche de la première plage.

2. Procédé selon la revendication 1, caractérisé en ce que la génération du filtre à partir du premier ensemble d'images acceptables du produit comprend la combinaison (120) du premier ensemble d'images acceptables en utilisant une opération logique OU.

3. Procédé selon la revendication 2, caractérisé par ailleurs en ce que l'étape de combinaison du premier ensemble d'images acceptables comprend les étapes consistant :

à détecter les bords (108) de chaque image; et à associer (108) une première valeur à chaque partie de chaque image dont les bords ont été détectés qui a une valeur sur un côté de la valeur de seuil prédéterminée, et à associer une seconde valeur à toutes les autres parties de chaque image dont les bords ont été détectés avant de combiner les images en utilisant l'opération logique OU.

4. Procédé selon la revendication 3, caractérisé en ce que la première et la seconde valeurs sont les valeurs binaires 0 et 1.

5. Procédé selon les revendications 1 à 4, caractérisé en ce que l'entraînement adaptatif du filtre est analogue à un algorithme d'entraînement de Widrow-Hoff.

6. Procédé selon l'une quelconque des revendications 1 à 5, caractérisé en ce que l'image sélectionnée est sélectionnée manuellement (324) sur la base de l'aspect du produit associé à l'image.

7. Procédé selon l'une quelconque des revendications 1 à 6, caractérisé par l'entraînement adaptatif du filtre avec l'image associée si l'autre valeur traitée se situe dans la première plage.

8. Procédé selon l'une quelconque des revendications précédentes, dans lequel les valeurs traitées de chaque image du second ensemble sont le produit scalaire du filtre et de chacune du second ensemble d'images.

9. Procédé selon l'une quelconque des revendications

précédentes, caractérisé en ce que l'on subdivise chacune des images acceptables en une pluralité de parties d'images et on effectue toutes les étapes précitées sur au moins certaines des parties d'images individuellement.

5

10. Appareil pour déterminer l'acceptabilité d'un produit (12) en générant un filtre (F) à partir d'un premier ensemble d'images acceptables (I) du produit, en comparant le filtre à chacune d'un deuxième ensemble d'images du produit pour produire une valeur traitée (P) pour chaque image du deuxième ensemble, les valeurs traitées ayant une distribution de valeurs, en comparant le filtre à chacune d'un troisième ensemble d'images du produit pour obtenir une autre valeur traitée pour chaque image du troisième ensemble, et en comparant chaque autre valeur traitée à la distribution pour déterminer l'acceptabilité du produit associé à l'image ayant l'autre valeur traitée, caractérisé par :

10

15

20

des moyens (26, 220) pour générer à partir de la distribution une première plage (A) et une seconde plage (B) de valeurs traitées, la première plage comprenant des valeurs traitées associées à des images acceptables du produit, et la deuxième plage étant espacée de la première plage et comprenant des valeurs traitées associées à des images inacceptables du produit; des moyens (26, 318, 322, 324) pour déterminer si chaque autre valeur traitée est extérieure à la fois à la première et à la seconde plages et, si c'est le cas, à sélectionner l'image associée uniquement si ladite image est acceptable; et

25

30

35

des moyens (26, 334) pour entraîner de manière adaptative le filtre avec l'image sélectionnée afin de produire un filtre modifié, de sorte que la comparaison du filtre modifié à l'image sélectionnée produise une valeur traitée modifiée plus proche de la première plage.

40

11. Appareil selon la revendication 10, caractérisé en ce que le filtre est généré à partir du premier ensemble d'images acceptables du produit par des moyens (26, 120) pour combiner le premier ensemble d'images acceptables en utilisant une opération logique OU.

45

12. Appareil selon la revendication 11, caractérisé en ce que les moyens pour combiner le premier ensemble d'images acceptables :

50

des moyens (26, 108) pour détecter les bords de chaque image; et
des moyens (26, 108) pour associer une première valeur à chaque partie de chaque image dont les bords ont été détectés qui a une valeur

55

sur un côté de la valeur de seuil prédéterminée, et à associer une seconde valeur à toutes les autres parties de chaque image dont les bords ont été détectés avant de combiner les images en utilisant l'opération logique OU.

13. Appareil selon la revendication 12, caractérisé en ce que la première et la seconde valeurs sont les valeurs binaires 0 et 1.

14. Appareil selon l'une quelconque des revendications 10 à 13, caractérisé en ce que les moyens (26, 334) pour l'entraînement adaptatif effectuent une fonction analogue à celle d'un algorithme d'entraînement de Widrow-Hoff.

15. Appareil selon l'une quelconque des revendications 10 à 14, caractérisé en ce que les moyens de détermination comprennent des moyens (26, 324) pour permettre une sélection manuelle basée sur l'aspect du produit associé à l'image.

16. Appareil selon l'une quelconque des revendications 10 à 15, caractérisé par des moyens (26) pour l'entraînement adaptatif du filtre avec une image associée à une autre valeur traitée qui se situe dans la première plage.

17. Appareil selon l'une quelconque des revendications 10 à 16, caractérisé par des moyens pour former le produit scalaire du filtre et de chacune du deuxième ensemble d'images pour former les valeurs traitées de chaque image dudit deuxième ensemble.

18. Appareil selon l'une quelconque des revendications 10 à 17, caractérisé par une caméra vidéo (24) pour former au moins l'un desdits premier, deuxième et troisième ensembles d'images et par des moyens pour positionner un produit dans le champ d'observation de la caméra vidéo.

19. Appareil selon la revendication 18, caractérisé par des moyens permettant d'éclairer le produit dans le champ d'observation de la caméra vidéo.

20. Appareil selon la revendication 19 ou 20, caractérisé en ce que les moyens de positionnement comprennent un transporteur (20) pour acheminer les produits l'un après l'autre à travers le champ d'observation de la caméra vidéo.




FIG. 1

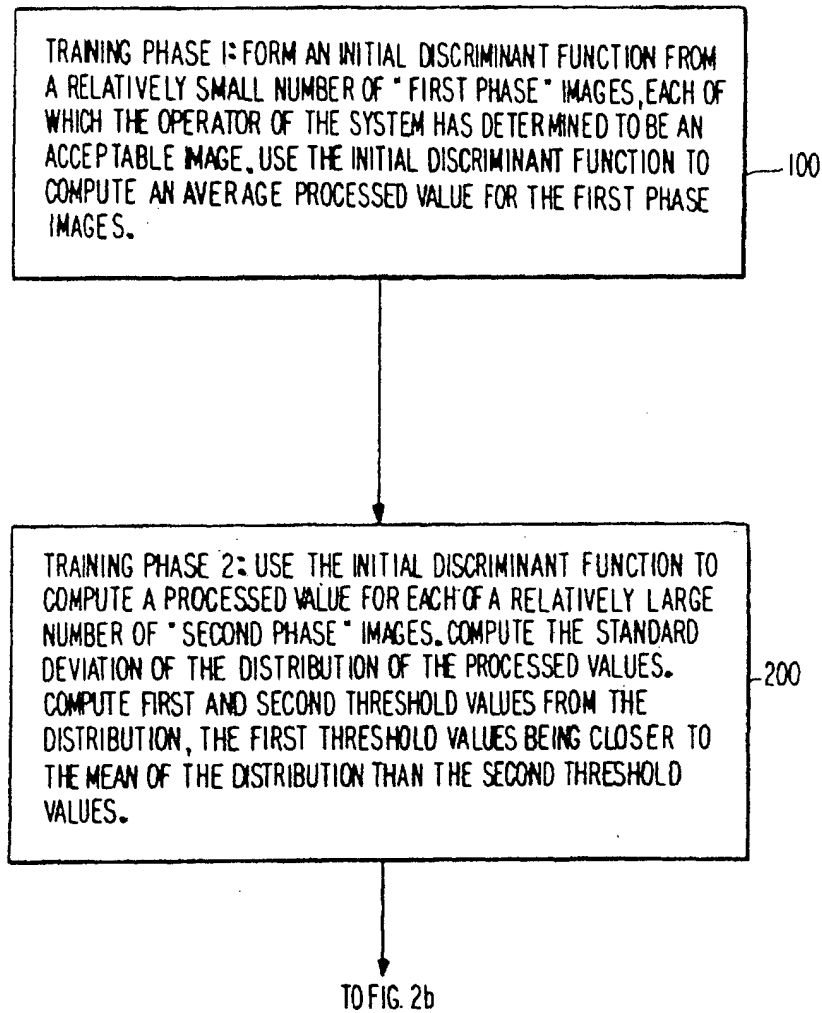
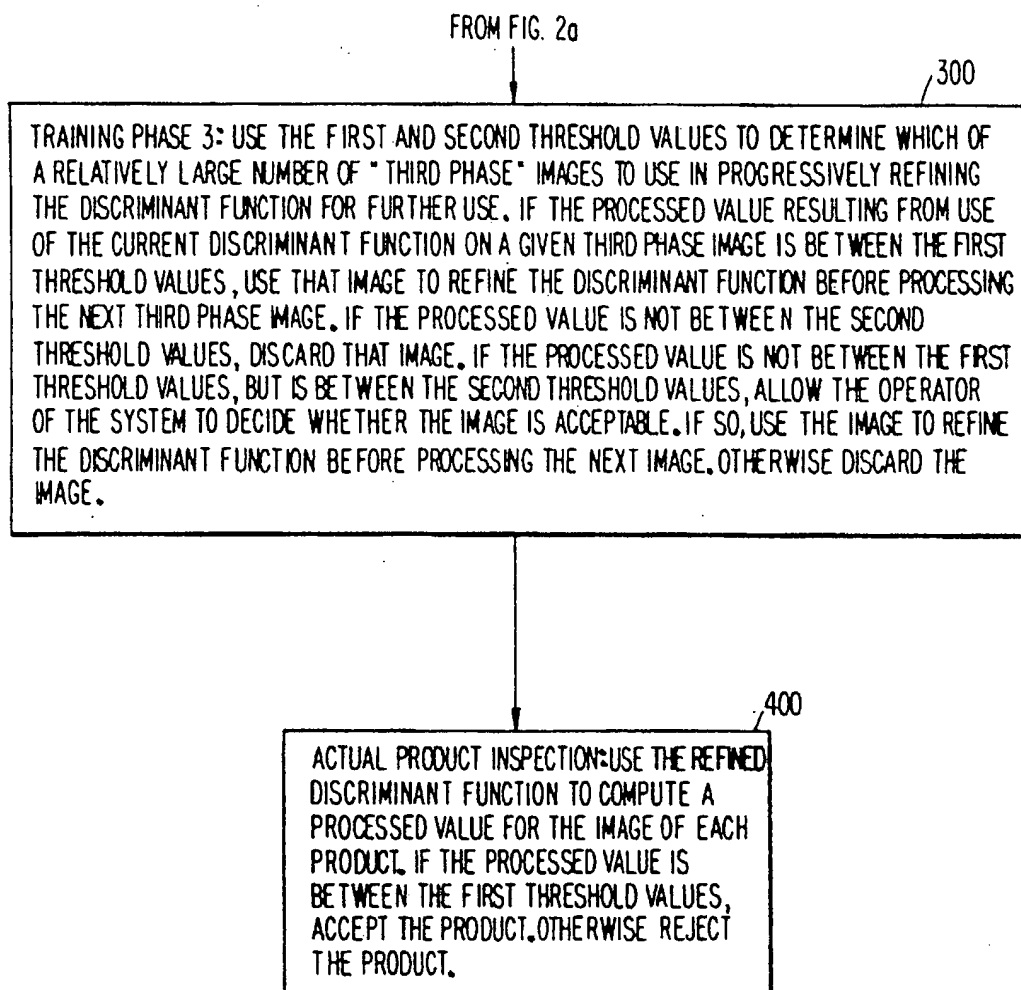
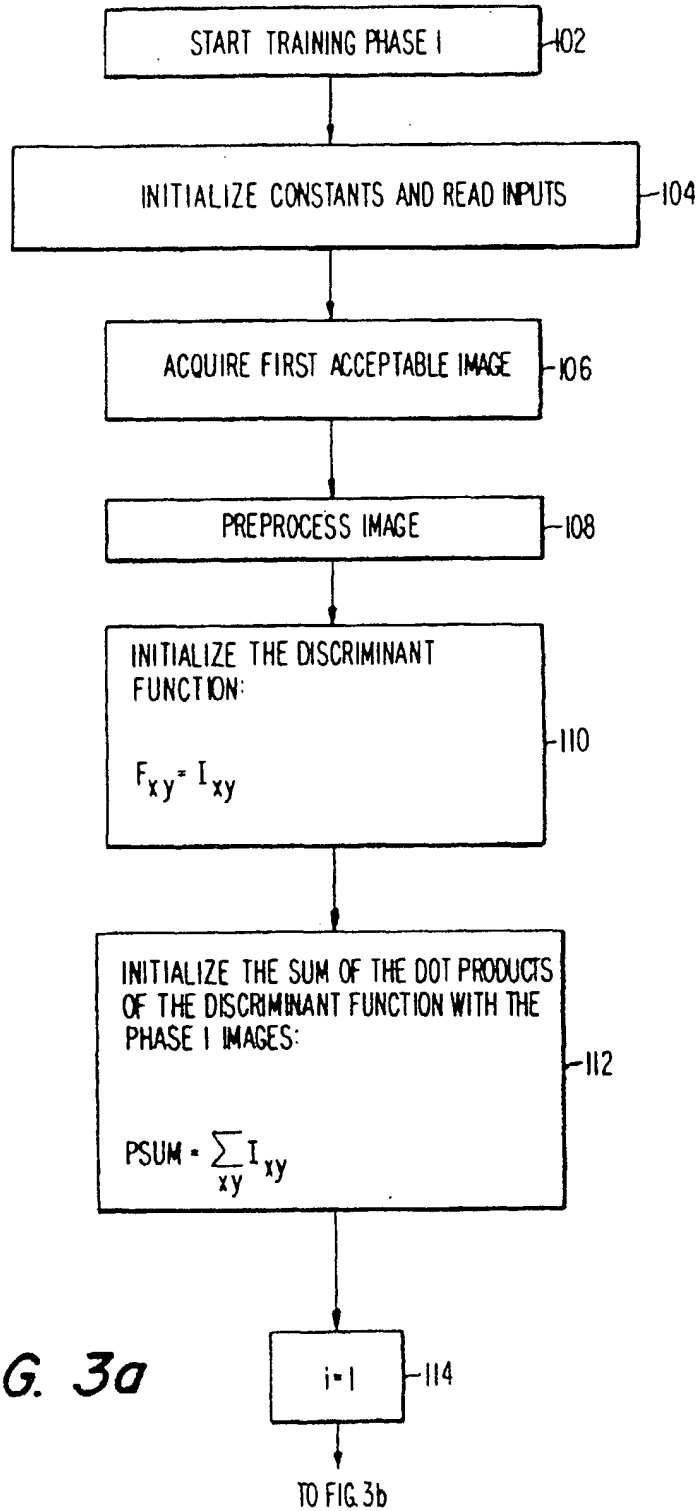


FIG. 2a

*FIG. 2b*



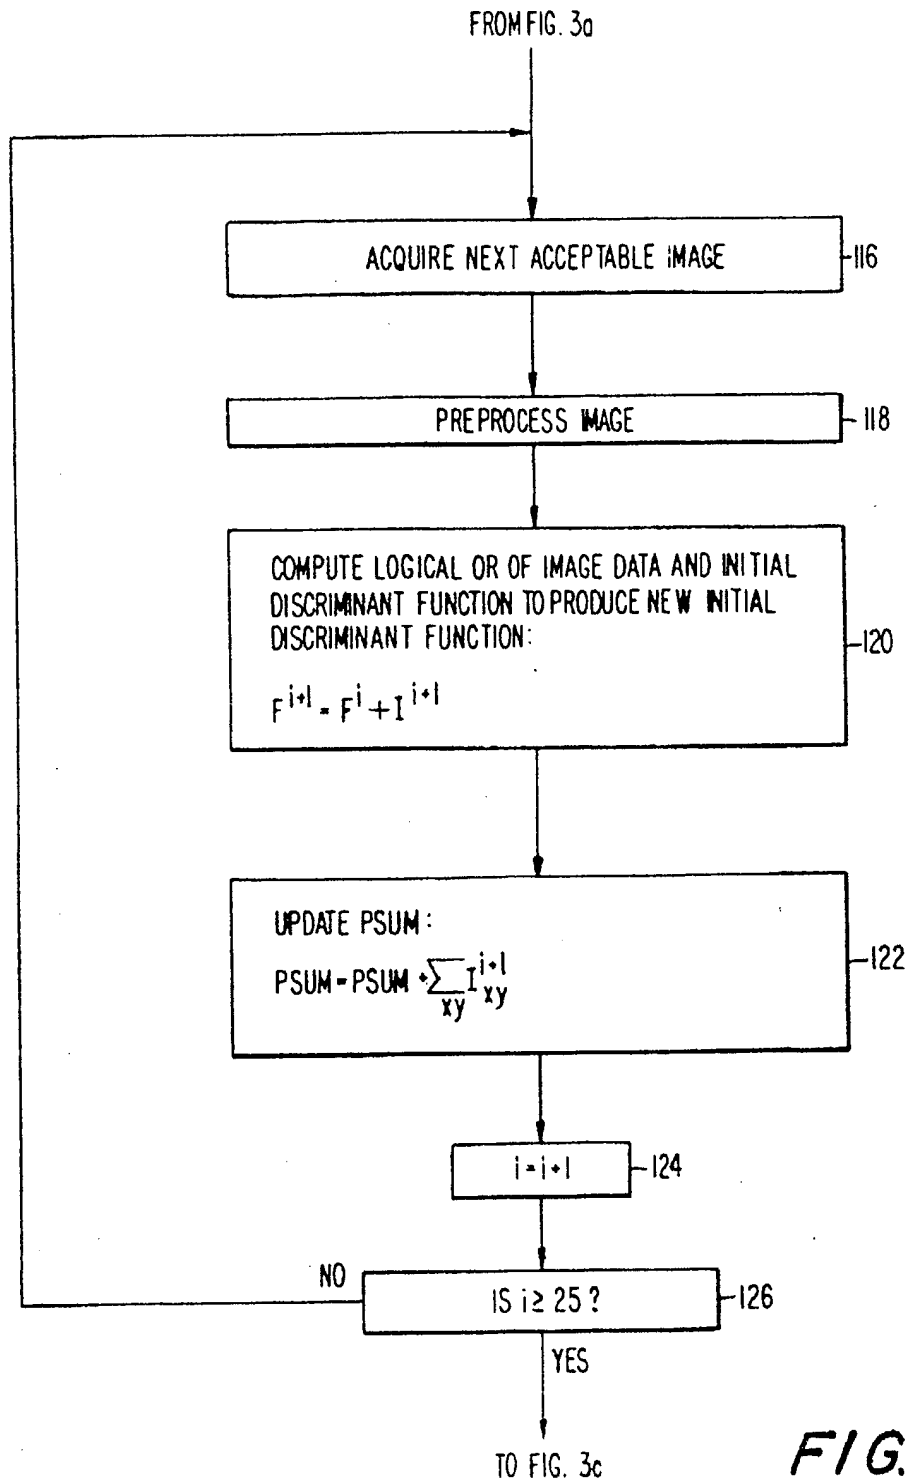


FIG. 3b


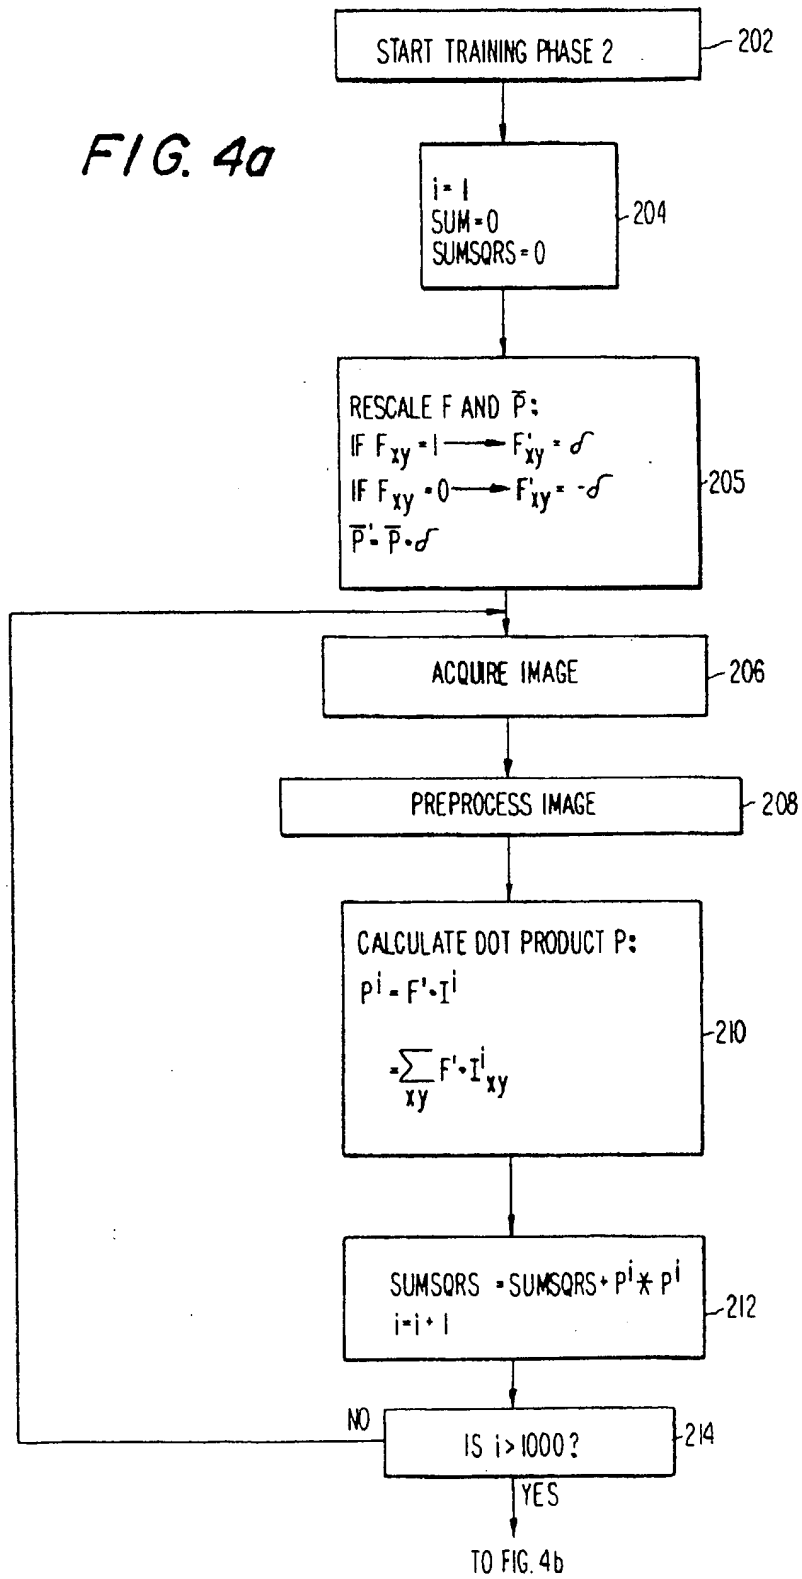


FIG. 3c

FIG. 4a



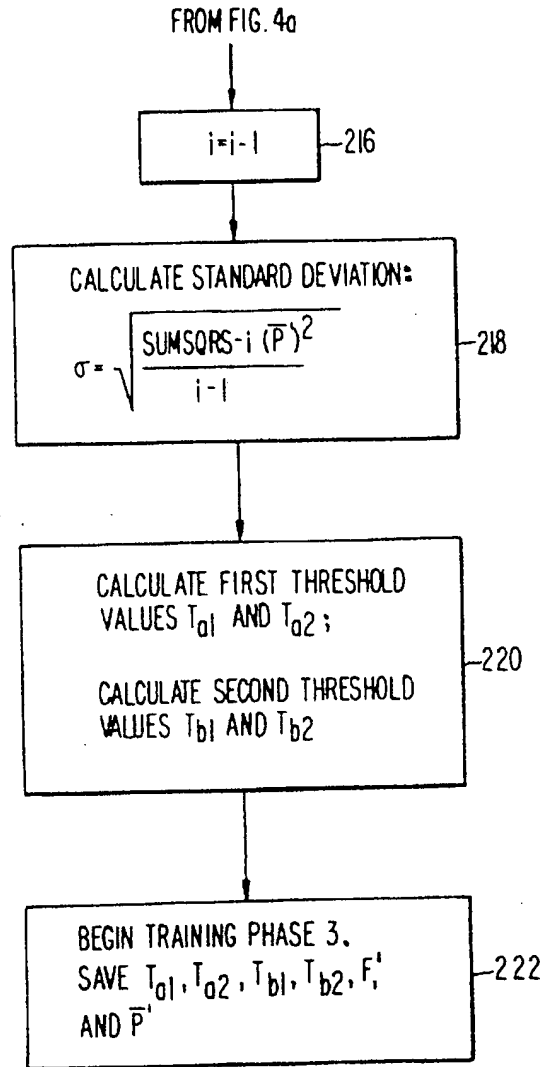


FIG. 4b




FIG. 5a

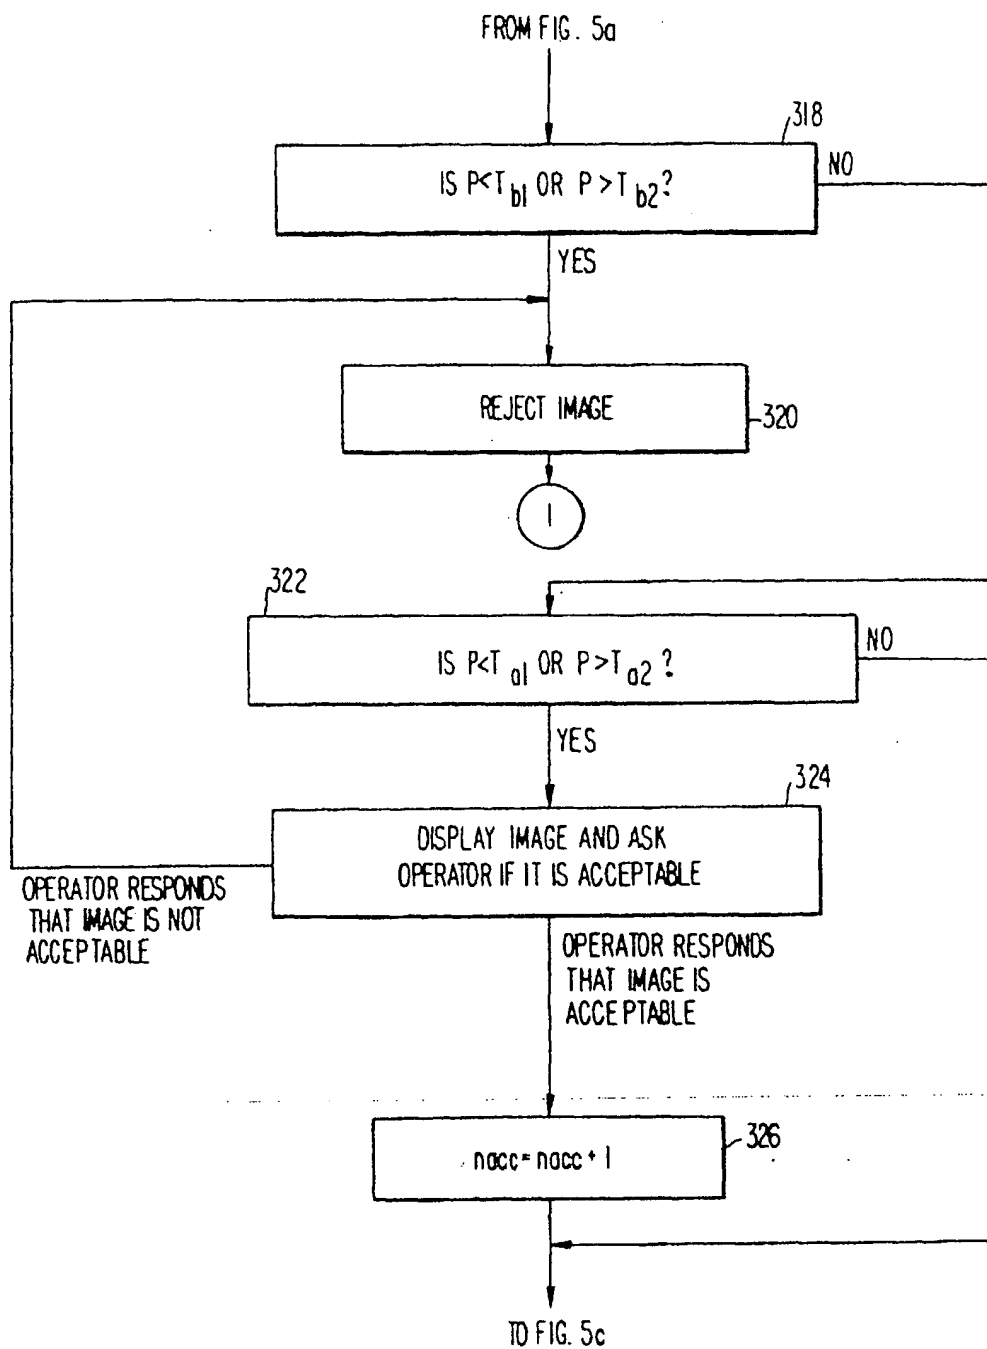


FIG. 5b

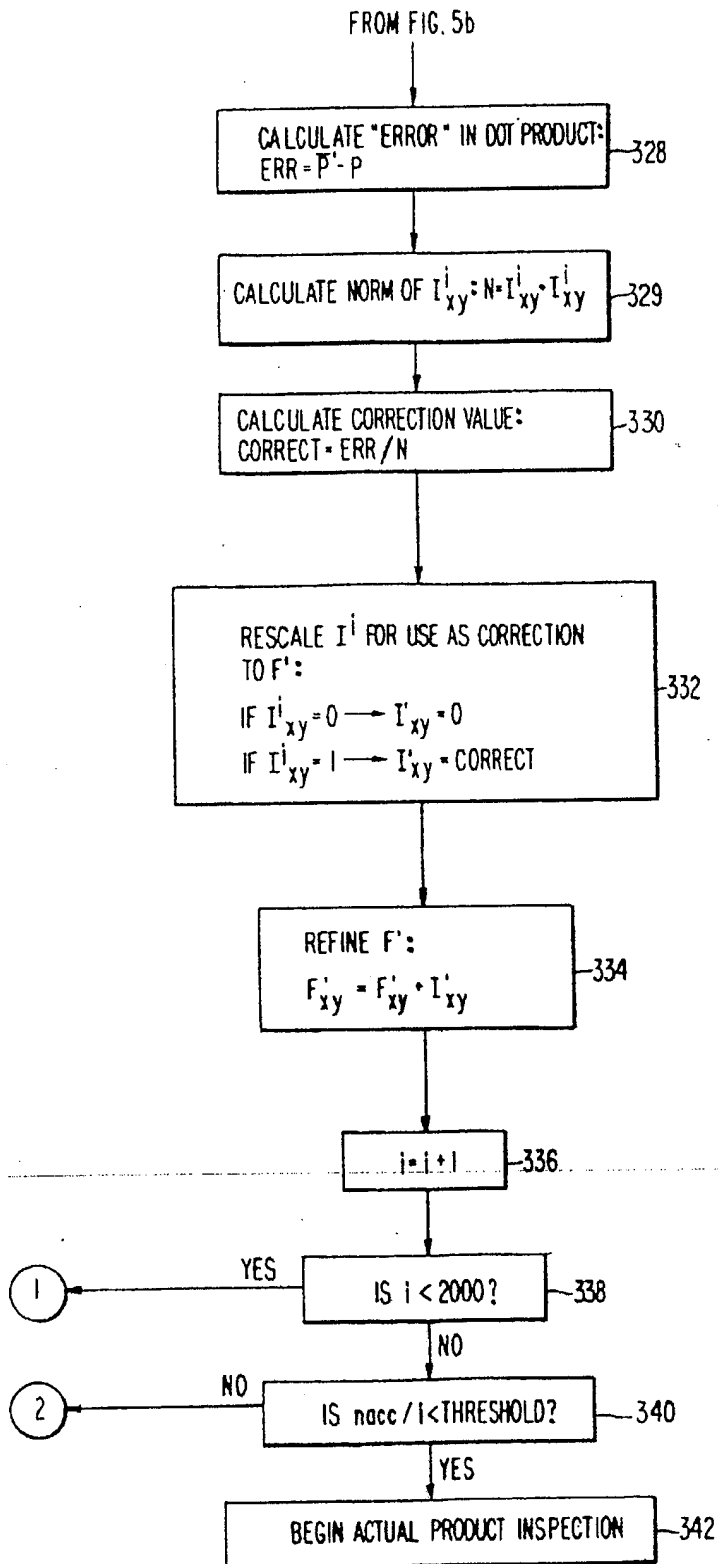
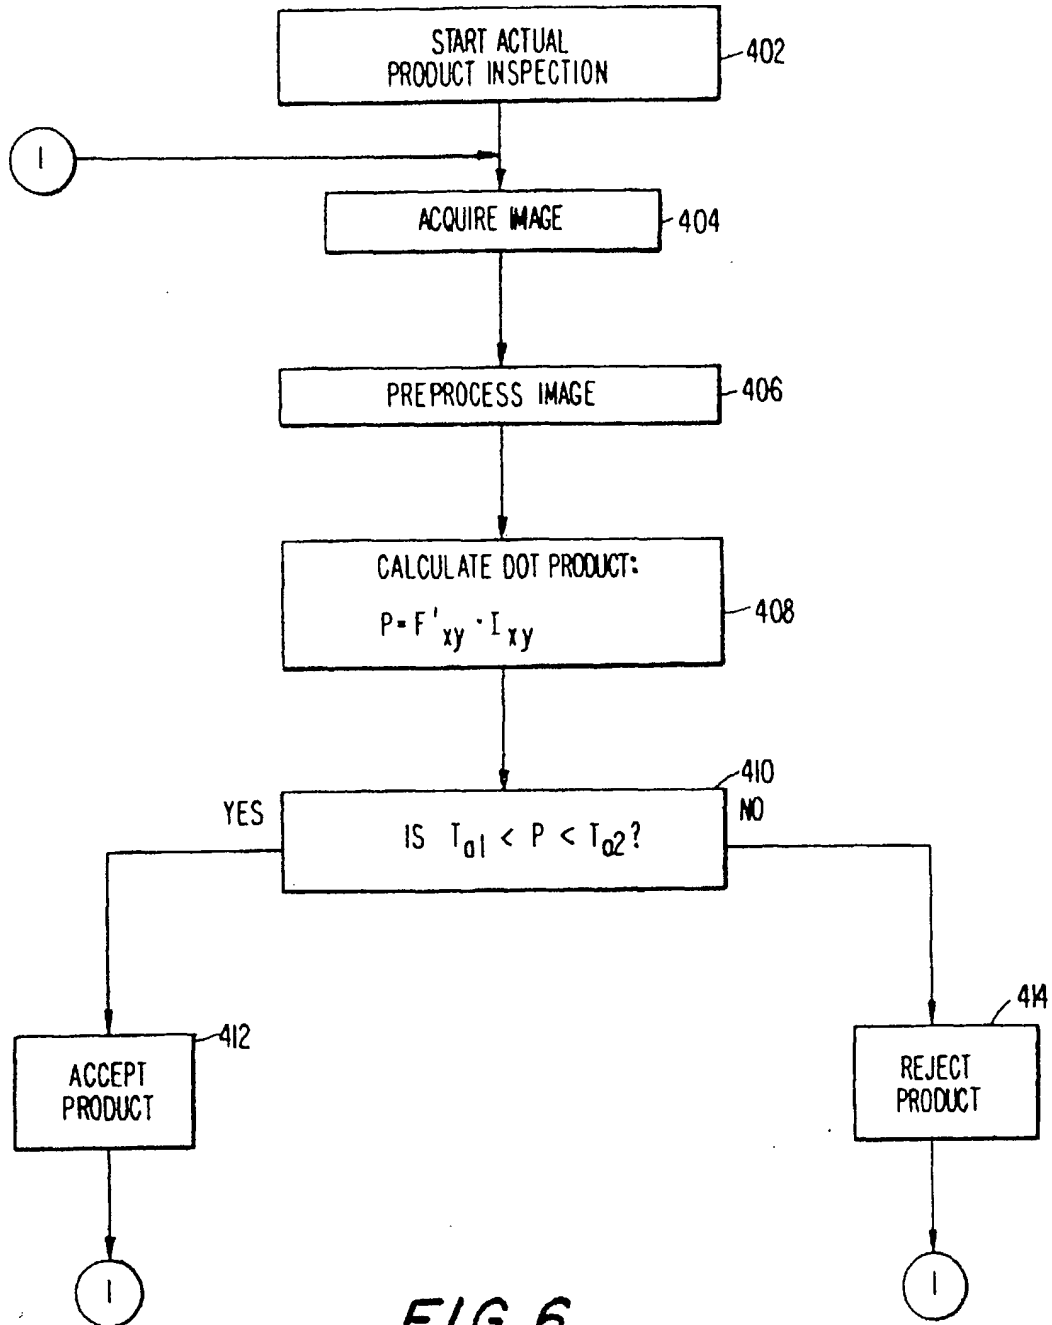


FIG. 5c






FIG. 7

FIG. 9

$$\begin{aligned} I_{d1} &= \bar{P}' - \alpha_d \sigma \\ I_{d2} &= \bar{P}' + \alpha_d \sigma \\ I_{b1} &= \bar{P}' - \alpha_b \sigma \\ I_{b2} &= \bar{P}' + \alpha_b \sigma \end{aligned}$$

FIG. 8

